

Main Injector Rookie Book

Chapter 3: Power Supplies

This chapter will discuss how electrical power is supplied to the magnets in the tunnel and to the equipment in the service buildings. It will also cover the regulation of the magnet power supplies, especially that of the major busses. The chapter is primarily concerned with power to the magnets in the ring; power to devices in the beam transport lines and to general service building utilities will be covered in later chapters.

The biggest consumers of electrical energy in the Main Injector are the large dipoles and quadrupoles. All of the dipoles are on a single bus, in series (Fig. 2-4). All of the focusing quads are on a second bus and the defocusing quads on a third (Fig. 2-7). The power delivered to these three sets of magnets is called pulsed power. Because of the large amounts of power required by these magnets, and the rapidly changing loads that they create, pulsed power is usually treated separately from other applications. The rest of the distributed power—to the smaller magnets, service building electronics, water systems, light bulbs, etc.—can be considered conventional or house power.

Commonwealth Edison and the Substations


Commonwealth Edison, or ComEd, is the public utility company that supplies power to most of the residential and commercial users in northern Illinois. Fermilab purchases all of its power from Com Ed. Fermilab's utter dependence on Com Ed becomes evident when a substation fire or an aggressive cherry tree interrupts power to the site, and sometimes weeks are required to recover from the damage. To be fair, Fermilab's rapidly changing inductive loads, especially the Tevatron and the Main Injector, can be a major headache for Com Ed as well.

ComEd generates power from several nuclear power plants distributed around northern Illinois. There are also coal-fired plants that can be called into service during times of heavy demand, such as on hot summer days

Main Injector

when 15 million air conditioners are operating simultaneously. Generated power is placed on the grid, a large network for the distribution of power. The transmission lines on the grid are usually maintained at 345 KV or higher. (Power lost as heat is proportional to current; since power is the product of voltage and current, a given amount of power is most efficiently transmitted with a high voltage and a low current.) The transmission lines entering the Fermilab site are at a voltage of 345 KV.

The 345 KV is 3-phase at 60 Hz; that is, there are three separate power lines, running parallel to each other, each being 120° out of phase with respect to the other two. The phases are usually designated “A”, “B”, and “C”:



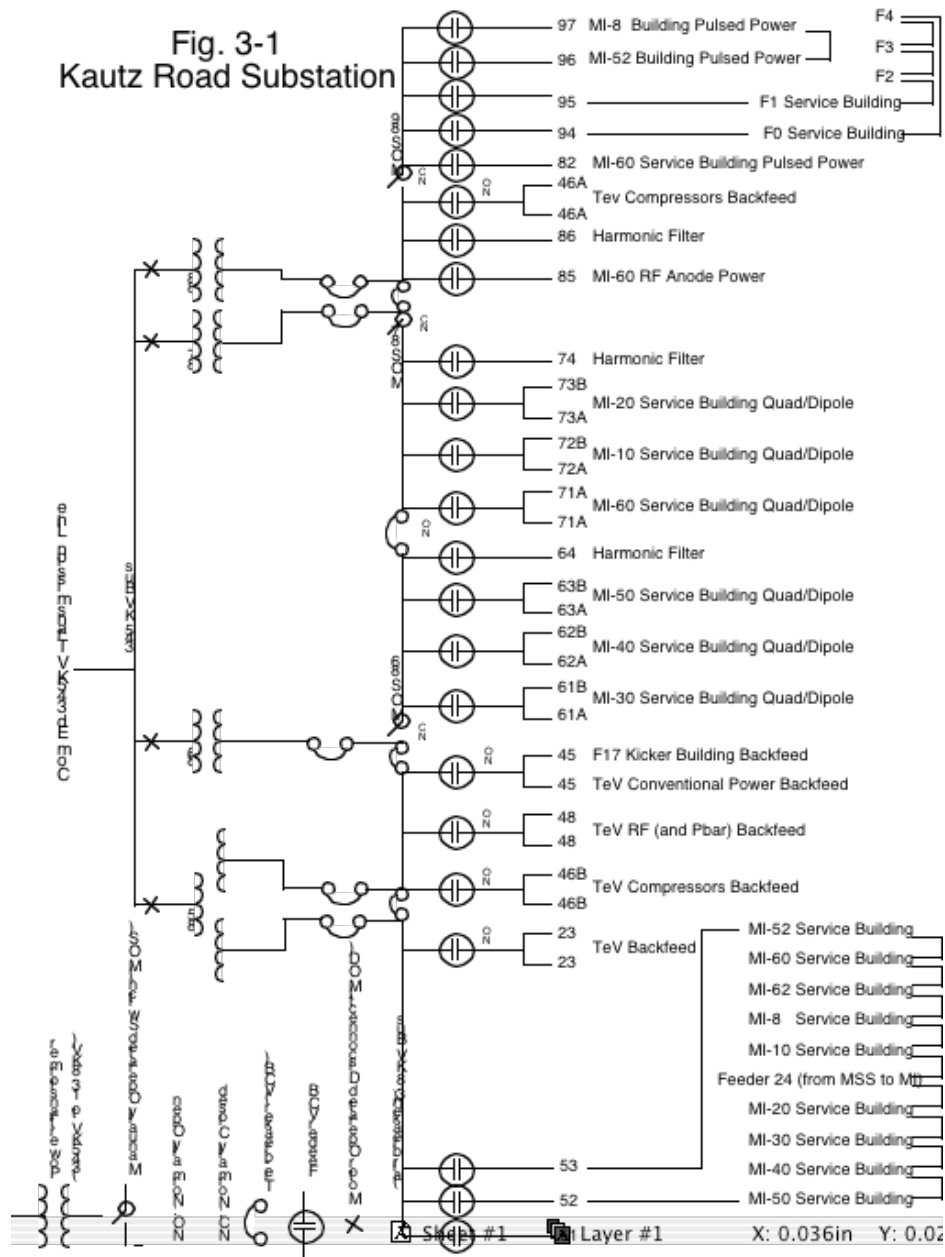
Wherever power is to be drawn from the grid, a substation is present. Transformers at the substations step down the 345 KV, a voltage generally too high for most household appliances, to a more manageable level. At Fermilab, the substation transformers step the voltage down to 13.8 KV.

There are three substations on the Fermilab site. The first is the Master Substation, located about half a mile north of the High Rise. The Master Substation supplies power to many of the accelerators as well as to most of the buildings on site. The second is Giese Rd., a small substation just to the west of the Antiproton Source. The Kautz Rd. Substation (KRS) is located just outside the Tevatron Ring near the MI-50 Service Building. It

Main Injector

was built specifically for the Main Injector and will be the one of greatest interest in this book. However, Giese Rd. and the Master Substation can “back-feed” some power to the Main Injector when necessary, and Kautz Rd. can back-feed power to systems normally fed through the Master Substation.

Fig. 3-1, located below, provides a general schematic of power distribution related to the Kautz Rd. Substation. It should be referred to frequently when reading the next few paragraphs.



Main Injector

The 345 KV lines to the Kautz Rd. Substation run parallel to Butterfield Road near the southern border of the laboratory until they strike northwest to meet the substation. Inside the substation fence the three phases connect to gigantic transformers, designated T-85, T-86, and T-88. They all step the 345 KV down to 13.8 KV. (The numerically alert will notice that T-87 is missing from the sequence. It was to be installed with the others, but it failed on the manufacturer's test stand and is still awaiting shipment. As of this writing (October 1999), it appears that the substation will be short one transformer for another year or so. T-88 has temporarily taken over much of the work that T-87 was to have done. The discussion to follow describes the current temporary configuration and will have to be revised when the replacement transformer arrives.)

From time to time, T-85, T-86, and T-88 need to be isolated for maintenance or repairs. On the primary (345 KV) side, air breakers, also known as Motor Operated Disconnects (MODs), at the top of each transformer, are used to isolate Com Ed's 345 KV from the transformers. A MOD operates by rotating a segment of the line and breaking the path of the current.

On the secondary (13.8 KV) side, vacuum circuit breakers (VCBs) are used to isolate the transformers from the load. VCBs are breakers that open and close in an evacuated "jar;" under vacuum, there is less likelihood of arcing. These main VCBs, rated to withstand 3000 amps, are also known as tiebreakers. (The analogy is obviously patterned after the duty of the Vice President to resolve a deadlock in the Senate.) There are also tiebreakers for isolating segments of the 13.8 KV bus.

The 13.8 KV from the transformers is distributed to the service buildings through individual feeders. (A feeder is a large underground cable capable of carrying a correspondingly large amount of power.) Unlike the Main Ring feeders, which had been buried directly in the ground and were susceptible to ground faulting, the Main Injector feeders are protected in underground concrete enclosures. They are often named after the last digit of the transformer from which they originate (e.g., feeders originating from T-86 are assigned numbers beginning with "6," and those from T-88 sometimes begin with an "8").

Main Injector

T-85 has two secondaries, one that supplies conventional power to the Main Injector service buildings and another that can be used for back-feed to the Master Substation. Feeders 52 and 53 carry conventional power from T-85 to the Main Injector ring. T-86 and T-87 were designed to provide pulsed power to the main dipoles and quadrupoles in the Main Injector ring; since T-88 has taken over T-87's load, many of its feeders begin with "7." The once and future purpose of T-88 is to supply pulsed power to beamlines, including the P1, P2, and P3 lines. Except for the P3 line, which has not yet been commissioned, T-88 still performs that function.

The individual feeders can also be isolated from their loads by VCBs. These feeder breakers, required to carry less current, are rated to 1200 amps. In addition to the tiebreakers, the 13.8 KV bus can be broken into segments by manually operated switches (MOS). These knife switches at Kautz Rd. are opened during an access. The reader should be able to determine from Fig. 3-1 that if all three MOS switches (MOS 86, MOS 87, and MOS 89) are open, pulsed power is removed from the bus while still leaving conventional power available to the buildings.

The procedure for opening the MOS switches is popularly known as racking out. In the Main Injector tunnel, the pulsed power is carried on exposed copper busses. Since this situation would represent a serious electrical safety hazard to anyone present, it is vital to confirm that the feeders cannot be powered. Each of the MOS switches has a dedicated "Open" and "Closed" key; the keys are controlled from the MCR. In addition, the VCBs for each pulsed power feeder are opened.

Remember that three-phase is present at every point discussed so far. Each MOS switch actually consists of three knife switches, one for each phase. A window has been provided for confirmation that all three knife switches have opened.

It should be noted that the MOS switches were not designed to be opened under load; all pulsed power must be removed from the appropriate devices before operating the switch.

Main Injector

Some of the feeders are paired and feed a system from opposite directions. For example, Feeder 52 starts its rounds at MI-50, while Feeder 53 starts at MI-52. The service a building disconnects can be configured so that the load is shifted from one feeder to the other.

Most of the feeders isolated by MOS 89 are designated with a “9.” (Remember, there is no Transformer 89; these feeders are powered from T-88.) Most of these feeders are dedicated to F Sector pulsed power, but 96 and 97 serve MI-8, MI-40, MI-52, and MI-62. Note that these locations are where most of the beam transport lines reside.

Pulsed Power

Remember that T-86 and T-88 provide all of the pulsed power for the main dipoles and quadrupoles. The feeders from T-86 travel underground from the substation until they reach MI-50. Two of the feeders, 63A and 63B, supply MI-50 with power. The other T-86 feeders turn to the left. 62A and 62B supply pulsed power to MI-40, and 61A and 61B continue on to MI-30. Similarly, the feeders from T-88 branch to the right and feed MI-60, MI-10, and MI-20. Feeders 71A and 71B go to MI-60, feeders 72A and 72B to MI-10, and feeders 73A and 73B to MI-20. When T-87 is installed, it will take over Feeders 71, 72, and 73.

Remember that in addition to the MOS switches, which isolate the feeders from the transformers, the individual feeder breakers are opened prior to an access. The breakers are operated from a console at the Kautz Rd. Substation.

Because of the transient spikes created when the feeder breakers are opened and closed, they are manipulated one at a time in a specific order. The order is executed by programmed logic controllers (PLCs), which act as software relays. PLCs have a variety of uses in the Main Injector; more detail will be forthcoming in the chapter on controls.


One feeder from each of the pulsed power transformers is connected to a harmonic filter. Each of the transformers has a dedicated harmonic filter;

Main Injector

the filters are all located in the northwest corner of the substation. Their purpose is to suppress frequencies near the 720 Hz and 1440 Hz (and other unwanted frequencies) that may be on the feeders. The 720 Hz and 1440 Hz noise, as will be explained later, originates with the power supplies in the service buildings and feeds backwards toward the pulsed power transformers. These high frequencies are added to the voltage already present on the feeders and increase the voltage stress on the system. The harmonic filters, which consist of large stacks of inductors and capacitors, greatly enhance the lifetime of the pulsed power transformers by screening off the higher frequencies.

It may seem odd that the filters are not directly in line with the feeders, but remember that the filters, magnets, and power supplies are all part of the same interconnected resonant system. The filter feeders (no, these aren't clams) are designed to work in parallel with the other feeders.

A large resonant system such as the Main Injector power supply/magnet network can generate side effects that would not be intuitive to the untrained observer. One of these is that a 200 Hz standing wave appears across the magnets. If not compensated for, the varying impedance would cause every magnet to have a current that is out of phase with that of its neighbors. The compensation comes in the form of a mode-damping resistor placed in parallel with each magnet:



Main Injector

Exactly how this resistor works to block the standing wave is beyond the scope of this author . . . I mean, book. That is, it is left as an exercise for the reader. However, it can be said that higher frequencies are blocked by the high inductance of the magnet but pass readily through the resistor, where some of the energy is dissipated as heat.

The Service Building Utility Yard

The term “Utility Yard” as used here refers to the area outside each service building where you find the transformers and other high voltage equipment.

The way in which pulsed power is processed in order to create current in the magnets can be exemplified by the supplies at MI-20 (refer to Fig. 3-2 for a schematic overview). The process is similar at all of the major service buildings.

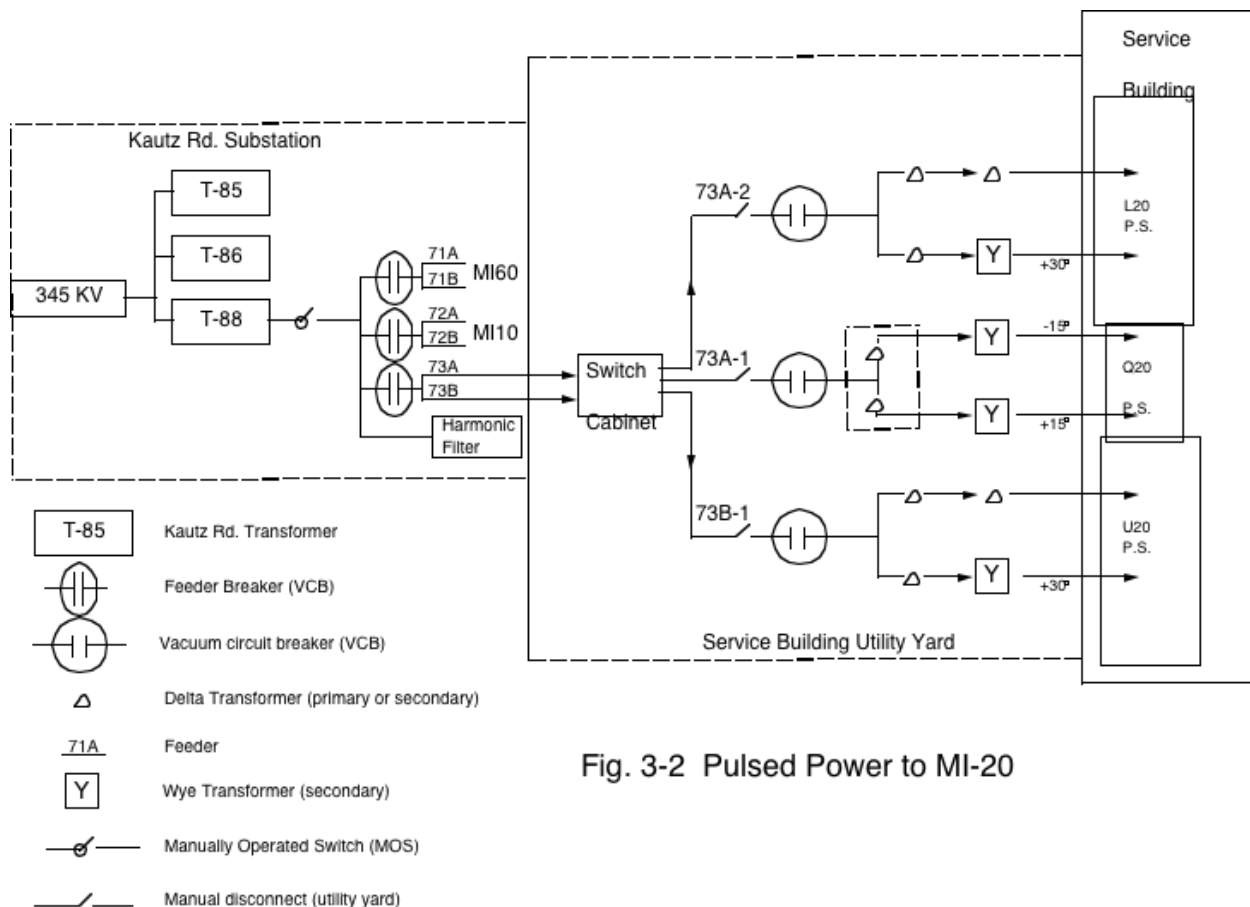


Fig. 3-2 Pulsed Power to MI-20

Main Injector

There are two dedicated pulsed power feeders leading to each service building. Remember that the feeders (73A and 73B) are energized to 13.8 KV. The first surface manifestation of power is the switch cabinet, a large box in the Service Building Utility Yard. Unlike the Tevatron and some other machines, these manual disconnect switches are not visible from the outside. These particular disconnects are used by FESS to reconfigure the feeders, not to isolate the power supplies.

From the switch cabinets the power is initially split three ways, one branch going to each bus (upper dipole, lower dipole, and quad). At this point there are two final opportunities to isolate the 13.8 KV from the power supply transformers. The first is the manual disconnect switch. This switch can be opened if work needs to be done on an individual supply. The second is a VCB, similar to those described earlier. The VCBs are opened routinely before an access or a hipot, and always when the permit loop is dropped. There will be more about those italicized items later in the chapter.

When the VCBs and manual disconnects are closed, power reaches the power supply transformers, which are also in the Utility Yard. The voltage output on the secondary windings of the transformers is about 1 KV. There is some further processing of the AC power by the transformers, but that discussion should be deferred until the power supplies themselves are described

There are six transformers altogether: two each for the upper and lower dipole supplies, and two for the quadrupole supply. (There are, of course, two quadrupole busses in the Main Injector ring, focusing and defocusing, but the power supplies for each bus are only found at every other building. At “even” numbered buildings such as MI-20 the quadrupole supply is dedicated to the focusing bus, and at “odd” numbered buildings it is connected to the defocusing bus.) The quadrupole transformers are both in a single tank, so it appears that there is only one.


Main Injector

AC to DC

The power supplies are located inside the service buildings. Their purpose is to convert the 60 Hz AC from the transformers to a DC voltage. This process is known as rectification (literally, “to make right”). The sine wave originating on the transformers spends half its time at the “correct” polarity and half its time at the “wrong” polarity. The rectification techniques involve (1) reversing the “incorrect” polarity and (2) packing the rectified sine waves so closely together that the output is nearly indistinguishable from a flat DC voltage. This approach is used many times over for power supplies at Fermilab and elsewhere.

The power supplies themselves are inside the service building. The core of the power supplies is a bank of silicon-controlled rectifiers (SCRs). An SCR is a device that, like a diode, allows current to pass in only one direction. An SCR has the additional feature of being triggered; without the trigger, the SCR will not conduct current even if the polarity is correct. If a conducting SCR is reverse biased, that is, if the voltage is reversed, it not only stops conducting but also has to be re-triggered before it can conduct again. Triggering, also known as gating, allows for precise control of the amount of voltage transmitted.

If a sinusoidal voltage is sent through an SCR (assume that it is gated), the waveform will look like this:



The SCR will only allow the positive portion of the waveform to pass. The dashed curve here indicates the portion of the waveform which was blocked; that is, where the voltage is negative. This type of rectifier is called a “half-wave” rectifier because only half of the available voltage survives the passage.

Main Injector

The specific type of power supply used for the dipoles is known as a full wave bridge rectifier (Fig. 3-3). In this simplified schematic, the transformer secondary is connected to a bridge consisting of four SCRs, and the load—in this case, a string of main dipoles or quadrupoles—is placed across the network.

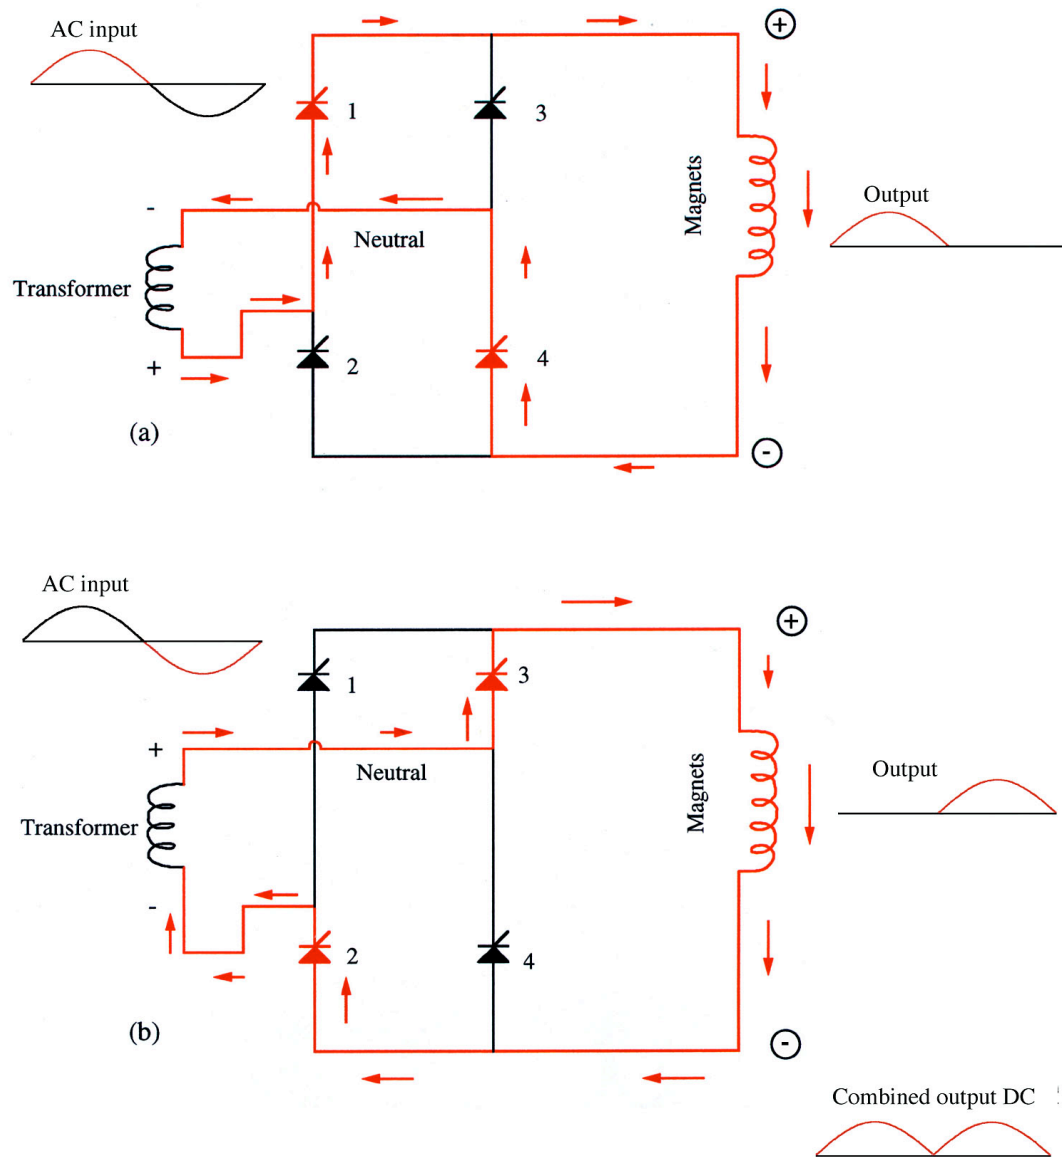



Fig. 3-3 Voltage Rectification using SCR's

This highly simplified diagram shows how the SCR's in the Main Injector power supplies convert an AC input to a DC output. See text for details.


Main Injector

Suppose that the current is to flow through the magnets as shown in Fig. 3-3(a). When the polarity of the voltage is positive (as in the picture above) the top SCR conducts and current flows through the magnets in the proper direction.


When the polarity reverses, as in Fig. 3-3(b), the bottom SCR conducts. The bridge is actually looking at the negative half of the waveform:



However, because of the way that the SCRs are connected, the direction of the current through the magnets is unchanged. The overall effect is to reverse the polarity of the negative portion of the wave. (The “full wave” designation comes from the fact that the entire wave is used.) The combined waveform is just beginning to resemble a steady voltage:



There is a new peak voltage every 180° , but remember that we are only looking at one phase out of three. When all three phases are considered (requiring more SCRs), there is a new peak every 60° :



The voltage is now being refreshed at a 360 Hz rate, which is not quite good enough, sort of like driving across a cobblestone road. But the transformers have one more trick up their proverbial sleeve, the primary coil


Main Injector

can be connected to two secondary coils instead of one, and the windings built so that the secondaries are shifted in phase with respect to each other. The dipole and quadrupole power supplies differ as to how this is done.

The two transformers on each dipole bus are of two different types. Both have a primary winding called a “delta,” named for the Greek letter that it resembles. The secondary windings are “delta” in one transformer and a “wye” in the other. (“Wye” is the letter “Y” spelled phonetically, for ease of pronunciation. It wasn’t my idea.) So there is one “delta-delta” and one “delta-wye” transformer for each dipole bus. The “delta” secondary retains the same phase as the primary, but the “wye” secondary shifts the phase by 30° . The phase shift doubles the number of sinusoidal peaks available to the power supplies, that is, it goes into the primary coil as 3-phase and comes out of the secondary as 6-phase.


The transformers and power supply components for the quadrupole power supplies have been recycled from old Main Ring supplies. Both transformers for a given supply utilize an “extended delta-wye” configuration. In one transformer, the phase of one of the secondaries is shifted forward by 15° , and in the other, the phase is shifted backwards by 15° . The resulting net shift of 30° is the same as what happens in the dipole supplies. The SCR network is similar, except that the individual SCRs are smaller. The chokes and capacitors in the filters are also less impressive than those in the dipole supplies, but the quadrupoles operate at a lower current, so they don’t have to be as impressive.

Combining the peaks coming from the three-phase input, the inverted negative portion of the wave, and the 30° phase shift, there are now 12 sinusoidal peaks per cycle. Voltage is being refreshed at a 720 Hz rate:



Main Injector

Of course, not all of the voltage is needed all the time; the SCRs are gated according to how much voltage is needed from the power supply at any given instant. The delay, measured along the sinusoidal wave, is called the firing angle:



The shaded area represents the time during which the SCR is conducting current. In isolation, the SCR would continue to conduct until the wave naturally reached a zero value. However, soon after the voltage begins to droop, an adjacent SCR fires and begins to conduct current from the next up-and-coming peak. Since the next peak is now at a higher voltage than the first, the first SCR is reverse biased and shuts off.

Remember that Fig. 3-3 is a simplified version of the actual SCR banks. In reality, there is a set of three pairs of SCRs for each of the three phases; then that number is doubled to account for the two secondaries feeding each supply, for 36 SCRs per supply.

This is as much resolution as can be achieved with the existing SCRs. More SCRs could always be added, but it would be impractical; most of the remaining 720 Hz ripple can be smoothed over with the chokes and capacitor banks inside the power supplies. A small amount remains.

The 720 Hz that enters the tunnel is blocked by the inductance of the magnets. A portion of the 720 Hz ripple created by the supplies, even though it is on the secondary side of the transformers, is picked up by the primary coils of the transformers in the utility yard. A 1440 Hz ripple (the first harmonic of 720 Hz) is generated as well. It is these two frequencies (technically, sidebands of those two frequencies), which find their way back

Main Injector

to the Kautz Rd. transformers, that the harmonic filters at the substation are designed to suppress.

There are also bypass SCRs associated with each power supply. The bypass SCRs shunts current away from the load. They are turned on routinely during every ramp cycle, depending on the specific requirements of the ramp at that instant. They are also activated if an unexpected failure in the power supplies or magnets requires that current be removed from the system. The fast bypass loop, described below, monitors conditions and fires the bypass SCRs if necessary.

The power supplies can be connected or disconnected from the magnet bus using the knife switches. These “large” knife switches can also be used to bypass the local power supply, leaving the bus unbroken so that the magnets can still be powered from the other supplies. A third configuration is to open up the knife switches completely, breaking the bus. This last option is used when searching for a ground fault.

There are “small” knife switches as well; they will be described shortly.

A signal from the potential transformer (PT) is also one of the inputs to the SCR firing module. A PT is not a transformer that has yet to live up to its ability, but rather a device for measuring the phase of the incoming voltage. All of the clever rectification techniques described above are useless if the SCRs don't know what the phase is.

The Power Supply Link

The voltage level produced by the power supplies depends entirely on the timing of the SCR firing triggers. If all of the voltage is blocked, there is no output; if none of it is blocked, the supply produces its maximum output. The local power supply hardware, given a voltage request, can calculate and implement the firing angles. The firing angles need to be updated with every new sine wave, at a 720 Hz rate.

In order to get all of the power supplies to produce just the right amount of voltage for the current needed in the magnets, the scheme must

Main Injector

be coordinated from a central location. A VME processor called MECAR, which resides at MI60N, transmits the voltage request to the power supplies via the power supply link (PSL). The link originates at MECAR; a cable launches the signal toward MI-10, counterclockwise around the ring. At each of the major service buildings, the signal is intercepted and decoded by CAMAC 269 cards, which then pass the information on to the local power supplies. These instructions are updated at 1440 Hz—twice as fast as the SCR firing angles are updated—with the hope that the increased resolution will carry over to 720 Hz. The local hardware knows how to convert the voltage requests into firing angles.

The only function of the CAMAC crate is to provide power and other basic services to the card; the link is independent of the PIOX, PIOR, and BTR links that CAMAC normally uses.

The health of the PSL link can be monitored from I17.

The scarcity of information about MECAR in this section will be rectified later in the chapter.

To the Magnets

Once the voltage has been produced, it must find its way to the magnets. The dipole busses emerge from the top of the power supply cabinets, surrounded by (but not touching) a metallic shield, and plunge downward into the tunnel near the door to the tunnel stairwell. They are bent horizontally when they reach the alcove downstairs.

The bus work is circular in cross-section when it leaves the power supply and enters the tunnel, but soon has to switch to the 1”X4” rectangle of the magnet coils (see Figs. 2-2, 2-3). The 1”X4” shape is generally retained when the bus has to pass behind the numerous quadrupoles, but reverts to a circular cross-section in the straight sections.

From a local perspective, it would seem that the upper and lower dipole supplies at each building are powering two separate busses, but in reality

Main Injector

the busses are continuous because they turn around at MI-60 (see Fig. 2-4, which the reader has surely memorized by now).

The Quadrupole Bus

An important difference between the dipole and quadrupole busses is that there is no fold in either of the quad busses. The focusing and defocusing busses are separate circuits, with current flowing one way in the focusing bus and the other in the defocusing bus. The individual busses require less current than the dipole bus, so there are only three power supplies per bus. The focusing bus is powered from MI-20, MI-40, and MI-60; the defocusing bus is powered from MI-10, MI-30, and MI-50.

MECAR controls the quadrupole busses in addition to the dipole busses.

The quadrupole busses enter the tunnel in a fashion similar to the dipoles, but their cross-section is always circular. Quads of a given polarity are connected with straight lengths of bus; there can be as many as three sections of quad bus running parallel to each other at a given location, and, because of the way they leapfrog, it is not always obvious which is which. Fig. 4-2 illustrates the quadrupole bus connections (in the context of LCW flow). One reason the buswork alternates is to accommodate the curvature of the tunnel; it is possible to use a straight length of bus in most locations. The other, and more important, reason is to allow the inductances and fields from each bus to cancel each other out.

The Permit and Fast Bypass Loops

The permit and fast bypass loops are used to ensure that the power supplies are not allowed to operate under adverse or unsafe conditions. When the permit loop is dropped, power is removed from the system by opening the VCBs; when the fast bypass loop is dropped, the bypass SCRs are turned on so that no current can be passed on to the load.

Main Injector

The permit loop is used for all situations involving personnel safety, because opening the VCB is the surest way to remove power from the system in a hurry. The Main Injector Electrical Safety System (ESS), the power supply door interlocks, and the big red “Panic Button” near each supply are all inputs to the permit loop. Certain overcurrent trips and “sudden pressure” indications are tied to the permit loop as well.

However, opening the VCB can cause transient voltage spikes if there is power flowing through it, possibly causing more damage. Situations that do not require such a drastic response use the fast bypass loop; these include problems such as high temperatures, water pressure, and transformer oil levels. In addition, the bypass loop is dropped whenever there is an unexpected ground fault; it is unwise to open the VCBs because of the voltage spikes just mentioned. The power supplies are in bypass whenever power is not available to the magnets, as when the 13.8 KV is racked out, or the VCB is open.

When turning on the Main Injector supplies, the permit loop must first be made up, and then the VCBs closed, before the bypass loop can be made up. Once the loops are made up, MECAR can be told, via the applications page I2, to begin sending current to the magnets.

The permit/fast bypass control chassis is located at MI-60 South. The chassis has buttons to push to bring the loops up, but fortunately, the applications page I17 provides a software interface to the MCR.

Fault Detection

The current traveling through the main dipole and quadrupole circuits must navigate through dozens of miles of copper busywork. Over the entire path, the bus must be adequately isolated from any conducting material that could cause a fault, including other parts of the bus itself. A turn-to-turn short inside a magnet would, for example, prevent part of the magnet from being powered, something which would obviously be detrimental to the beam trajectory. A ground fault will draw current through components, and two

Main Injector

simultaneous ground faults could draw enough power through the magnets to destroy them.

Faults can be generated in several ways: a coil-to-coil short can be caused by a breach of the magnet insulation, or a piece of copper tubing can be left wrapped around two busses after a maintenance day. In the Main Ring, high voltages on the bus would slowly cause copper to be electroplated inside the ceramic insulators that were supposed to isolate the LCW plumbing from the bus. It has been recently discovered in the Main Injector that copper is slowly absorbed by the thermoplastic hoses isolating the quadrupoles from the LCW headers, gradually creating a path for current flow from the bus to the grounded headers. Whatever the cause, it is important to continuously monitor the integrity of the three main busses in the Main Injector.

Ground faults and bus-to-bus shorts are detected through ground fault circuitry that is interlocked to the power supplies. In addition, a procedure called hipotting is employed, particularly following an access, to specifically look for potential faults. All of the main dipoles and quadrupoles are connected to a ground fault/hipot loop, which makes the hipot measurement straightforward.

Coil-to-coil shorts can be found by measuring the inductance across the magnets. This is a tedious and time-consuming process, since the magnets must be measured one at a time, and is not done as routinely as hipotting.

First, ground faults and hipotting.

The Ground Fault/Hipot Loop


The ground fault/hipot loop is used to search for faults on the Main Injector bus. It actually consists of two separate entities sharing the same cables. The cables run from building to building, surfacing in the knife switch cabinets where the magnet bus is connected to the power supplies. There are loops for the QF, QD, and Upper and Lower bend busses; they can

Main Injector

be dealt with separately, or in combination. In the knife switch cabinets, the cables are connected to the bus through a 50K resistor and a small knife switch (in contrast to the large knife switches that connect the power supplies to the magnet bus.) The resistor limits the amount of current flowing through the bus in the event of a ground fault. There is resistor and switch on each side of each power supply. Having many different paths to ground—in this case, through numerous current-limiting resistors—is known as a distributed ground system.

The hipot power supply is located in the back racks of the Main Control Room. The hipot controller itself is located at MI-60 South, near the controller for the permit and fast bypass loops. The software interface with the controller can be found on page I17.

Fig. 3-4 shows the hipot cabling for the upper and lower bend busses. This picture should be compared to Fig. 2-4; also be aware that there is a similar setup for the quadrupole busses.



The Hipot loop consists of cables connected to the magnets and their power supplies through 50K resistors. There are voltage taps on the bus side of the resistors. The Hipot power supply places -1000V on the loop. If there is no ground fault, no significant current flows through the resistors, and the voltage taps all see the same voltage as that produced by the Hipot supply.

The Main Dipole Hipot Loop, which includes the upper and lower busses, is shown. There are also Hipot loops for each of the two main quadrupole busses.

Main Injector

All four busses share the ground fault/hipot electronics.


During normal operation, any current flowing to ground will return to the ground fault circuitry (via ground, of course) and will be measured by an ammeter. The ground fault detector circuit is used to catch faults if they occur when the machine is running. If a fault is found, the fast bypass loop is pulled.

To actively hunt for a ground fault, a switch connects the loop to the hipot power supply (the VCBs must be opened before making the switch). The hipotter provides a static voltage (up to 3 KV, although 1 KV is the norm) to the hipot cable. The hipotter is usually first used in the “Bus and Power Supply” configuration. In this configuration, the bypass SCRs are gated. If there is no ground fault, there will still be a small amount of current through the 50K resistors, but no appreciable voltage drop. (The more current flowing through a resistor, the greater the voltage drop.) Voltage taps on the bus to either side of the power supplies measure the voltage on the bus. If there is no voltage drop, the bus will be at the same voltage as the hipot loop. The voltage from the taps can be graphically displayed from page I17.

If, however, the bus is ground faulted at some point, current will be drawn through the resistors, and the voltage drop will be noted on the taps. Since the resistors are in parallel, a hard ground fault will mean that none of the taps will read any voltage.

Main Injector

If that happens, the next step is to go to the “Bus Only” configuration (Fig. 3-5 (a) and (b)). In this state, the bypass SCRs are not gated; a ground fault on one side of the power supply will pull down the voltage on that side, but not on the other. Since the situation is similar at the adjacent power supplies, the ground fault is now isolated to a region between the two power supplies.



A ground fault has developed on the upper bus. Current is flowing to ground through the 50K resistors to either side. The hipotter measures the current. In this case, enough current is flowing to cause all of the voltage to be dropped across the two resistors, as shown in the voltage plot at right.

This isolation technique may not work, however, if one of the transformer secondaries or some of the SCRs are shorted. To completely isolate the sector, both the large and small knife switches must be opened.

The other two hipot options look at coil-to-coil faults. In the P2 mode, the upper bus is grounded and the lower bus hipotted; any short between the busses is interpreted as ground current. In the P3 mode, the lower bus is grounded and upper bus hipotted. A bad result could mean that the busses are touching, or (worse) one of the magnets is internally shorted.

Main Injector

Gadgets for checking the direction of ground current in the tunnel include Groundhog and the Blacklister. Groundhog operates with a 45V supply while the Blacklister places up to 1 KV on the bus.

Inductance Measurements

Occasionally a magnet will develop a coil-to-coil short that does not produce a ground fault. If only a small part of a main dipole or quadrupole is shorted, the effect on the beam may be subtle enough for the cause to go undiscovered for months; much troubleshooting of poor beam quality may occur without finding the cause. Each type of main dipole or main quadrupole has a predicted inductance based on the geometry of its coils. Inductance probes, which have to be manually placed on the magnets one at a time, can measure the inductance. If the inductance of a magnet falls short of the ideal, it may mean that a short has disabled part of the coil, and the magnet will have to be replaced.

An inductance probe works by generating a high frequency signal and measuring how it is propagated through the magnet. A switch, usually located behind the magnet in the most awkward location imaginable, opens the mode-damping resistor so that it does not interfere with the measurement by attenuating the higher frequency components of the signal.

MECAR

MECAR stands for “Main Injector Excitation Controller And Regulator” (the superfluous words are included to make the acronym easier to pronounce, as Mee-kar). MECARs job is to orchestrate the current in the main dipoles and main quadrupoles; it physically inhabits a VME crate at the north end of MI-60. (There is a spare MECAR in the same rack. One is known as MECAR A, the other as MECAR B.) The software in MECAR controls the amount of current the magnets will see by regulating the voltage produced by the power supplies. Some of the issues are discussed in the following sections.

Main Injector

Momentum, Current, Voltage

Current in the dipoles is frequently expressed in terms of beam momentum rather than amperes, because the magnetic field is more or less proportional to current, and the beam has to be within a narrow range of momentum if it is to be constrained by a given magnetic field. The two concepts diverge somewhat in the Main Injector because the saturation of the magnets at high currents distorts the proportionality between current and momentum.

A note to purists about the following discussion: ask your average Joe standing in the grocery line about the relativistic properties of subatomic particles, and he will probably express his feelings in terms of energy, as measured in electron volts, without giving it a second thought. The MECAR tables, however, list the momentum of the beam in GeV/c. As mentioned in Chapter 2, relativistic momentum is the kinetic energy, in electron volts, plus the rest mass of the particle, which is also measured in electron volts. The rest mass of a proton or antiproton is about 938 MeV. MECAR lists the momentum at injection as 8.889 GeV/c. That is the same as a kinetic energy of 8 GeV. (All right, the numbers don't quite add up. This business of $\text{Momentum} = \text{Kinetic Energy} + \text{Rest Mass}$ is really just an old folk tale, but one based on fact. It becomes increasingly accurate with increasing energy. 8 GeV just isn't quite relativistic enough to hide the difference.) Anyway, the units are almost arbitrary. Before the request for momentum can be put on the power supply link, momentum has to be converted to magnetic field strength, the field strength has to be converted to the equivalent current in the magnets (adjusted for hysteresis and saturation), current has to be converted to voltage (adjusted to compensate for the inductance of the magnets), and voltage has to be converted to a complex set of SCR firing angles (subject to a resolution of 720 Hz). In an attempt to mimic everyday banter, this text will carelessly bounce back and forth between energy and momentum, but the reader should know that there is a difference.

Main Injector

The applications page I2 is the user interface to MECAR. I2 includes time vs. momentum tables, which dictate the amount of current in the main dipoles, as well as time vs. tune tables, which ultimately determine the current in the main quadrupoles. As a convenience to the tuner, I2 also includes sextupole and octupole tables, but, as will be seen later in the chapter, these two types of magnets depend on CAMAC 453 cards rather than MECAR.

The various calculations are divided between I2's software and MECAR. First, I2 figures out what the magnet current needs to be, based on the desired beam momentum. Detailed conversion tables are then consulted that consider saturation. (The relation between momentum and current is almost proportional up to 120 GeV or so, after which the increase of the field slows significantly compared to the increase in current.) Some compensation is made for hysteresis as well, making certain assumptions about previous ramps. When the calculations are complete, I2 sends a table of desired current to MECAR.

It is MECAR's responsibility to calculate exactly what voltage is needed to produce the desired current. To do that, it needs to know the inductance of each segment of the ring between the power supplies. The di/dt can be very high, up to 10,000 amps/second, so the inductive impedance can be very large during those times. Up to 750 volts per power supply may be required to keep the ramp going.

Page I2 has an additional step to complete for the quadrupoles. The tunes are listed in the tables. The gradient field establishes the tune, which is measured in, say, kilogauss per meter. However, the beam momentum is changing during the ramp; the gradient has to change to match the beam momentum just to keep the tune constant during the cycle. So, the quadrupoles have "baseline" ramps to match the momentum. I2, when calculating the current, increases or decreases the current request based on whether the requested tune is above or below the baseline curve.

Main Injector

Regulation

The calculations of I2 and MECAR get the current close to where it should be, but reality soon takes its toll. In such a large system there are hundreds of variables, temperature being the most important, that can influence the final value of the current. There are several mechanisms available for controlling the outcome:

- The individual power supplies, like all decent power supplies, have internal feedback to ensure that they are producing voltage as requested. They compare their output with the voltage request arriving from the power supply link, and adjust the SCR firing angles appropriately.
- MECAR monitors the current in the busses and adjusts the voltage program accordingly, in real time. There are four transducers at MI-60: one for each quad bus, and one each for the upper and lower bend busses. (Only one of the bend bus transducers is used at any given time, since they are in series with each other.) MECAR reads the current on the three busses and implements a fast feedback loop to adjust the voltage as necessary. Any one of the 12 bend supplies can be selected as the bend bus regulator. Each of the quad busses also has a designated regulator. The fast feedback information is sent only to the regulators; the remaining supplies continue to play the basic pre-calculated waveform.
- If a ramp of a given type continues to deviate from its expected output during a cycle, MECAR remembers the error and applies a correction to the next cycle. This is called learning. These updates are semi-permanent and reduce the need for fast feedback. Updated voltage information is sent to all of the power supplies, not just the regulators.

Main Injector

The overall voltage request for a given cycle is called a profile. By “saving” a profile, the learning updates are assimilated by the profile, although they will be lost if MECAR is rebooted. Sometimes, however, the profiles are deliberately cleared of their updates. This is often done after MECAR learns in garbage (a universally accepted scientific term) when the ramp is faced with less-than-optimal conditions.

Ramps

The discussion in this section will emphasize the dipole (i.e., momentum) ramps. As mentioned earlier, the quadrupole ramps are generally proportional to the momentum ramps, with adjustments being made for changing tunes.

There are 8 GeV “ramps” where the beam momentum doesn’t change, but all the other ramps share some major milestones that occur during each cycle.

Injection always takes place at 8 GeV. Occasionally the amount of time that beam is circulating at 8 GeV is referred to as the dwell time.

During parabola, the current in the dipoles begins to increase. As implied by the name, the increase starts out slowly, the change in current (to a first approximation) literally following a parabola. As the current increases, the RF acceleration systems are changing the beam energy to match the strength of the magnets.

Near the top of the ramp, an inverted parabola eases the transition to flattop. The current in the magnets, and therefore the beam energy, does not change during flattop. The two flattop values used for Main Injector are 120 GeV and 150 GeV. During flattop, beam is at its highest energy for the cycle and is usually extracted from the Main Injector. Flattop is also the time when some of the more complex RF manipulations take place, such as bunch rotation and coalescing, but even if nothing special is done, the RF must still be present at some level in order to maintain the bunch structure of the beam and to compensate for energy lost to synchrotron radiation.

Main Injector

(Synchrotron radiation is a relativistic phenomenon and will be discussed in the RF chapter.)

After flattop, the ramp is driven into invert. The current in the magnets is taken back down to the 8 GeV value, and then some. Just because the energy is decreasing does not mean that invert is a passive process; the considerable inductance of the magnets would cause the current to linger for an unacceptably long time. It is not enough to shut off the voltage, the power supply voltage must actually be reversed during invert to force current out of the magnets.

The magnet current during invert is driven below the 8 GeV level in order to “reset” the hysteresis, ensuring that every cycle will begin with the same magnetizing force.

Finally, the magnet current bounces back to the 8 GeV level and the next cycle begins.

The total amount of time required for a complete cycle varies from about 2 seconds, in the case of antiproton production, to several seconds when injecting coalesced bunches into the Tevatron. Each type of ramp is designed to complete all of its necessary tasks as quickly as possible.

MECAR, then, must control three different ramps: those of the main dipoles, the main focusing quads, and the main defocusing quads. After making all of the appropriate calculations, the commands are sent to the power supplies via the power supply link. All phases of the ramp cycle (dwell time, parabola, linear ramp, flattop, and invert) are implemented at the power supplies by the timing of the SCR triggers.

The tables for the ramps can be found on the MECAR page, I2. There is a separate operational table for each type of reset (\$20, \$21, \$29, \$2A, \$2B, \$2D, or \$2E). Backing every operational table is a set of files. The tables determine the ramp in terms of time vs. momentum.

It is worthwhile to analyze a specific ramp using numbers from an actual MECAR table. The following discussion follows the fate of the main dipole magnet current over the course of a cycle. The \$21 cycle, which

Main Injector

accelerates five or six batches to 150 GeV, can be used as an example. The numbers in this example have been taken from operational tables, so they can be considered realistic, but not immutable, since they tend to be adjusted over time. Remember that these tables are reset-dependant; for example, a \$29 ramp will include many of the same basic features as the \$21 but will differ in many of the details.

The ramp has a dwell time of .55 seconds, during which the beam momentum is held constant at 8.889 GeV/c. That is sufficient time for all of the batches to be accelerated in the Linac and Booster, and to establish a stable orbit in the Main Injector. On I2, this line is labeled INITI, for “initial.”

Remember that there are twelve dipole power supplies altogether. At 8 GeV, only the regulator is producing current.

At .55 seconds, the current begins to slowly increase. Between 0.55 seconds and 0.5855 seconds, the momentum changes from 8.889 GeV/c to 8.96 GeV/c. This is the first part of the parabola, and of course, the rate of change in a parabola is slowest in the beginning and speeds up over time. This line is labeled VPARAB.

Notice that there are three numbers to the right of the “momentum” column. These are the derivatives, or rates of change, of the momentum. “Pdot” is the first derivative, or the rate of change of momentum. Pdot is measured in units of GeV/c per second if looking at momentum, or amps per second if looking at current. At .5855 seconds, the rate of change is 6 GeV/c per second.

To the right of the “Pdot” column are the second and third derivatives of momentum, Pddot and Pdddot. When Pdddot has a nonzero value, it means either that the ramp waveform equation has added a cubic term, or that there is a problem with the “d” stroke on the keyboard.

The regulator is the first supply in the turn-on order. The turn-on order—the sequence in which the supplies are turned on—is flexible and can be configured from Page I2. As the demand for current increases, the

Main Injector

regulator needs to be supplemented by some of the other supplies. The supplies are organized into tiers, with each tier comprised of a group of power supplies. Like the turn-on order, the tier structure can be modified, but it is important to distribute the supplies evenly around the ring to prevent the voltage to ground from becoming too high at any given point. For example, if the first tier consists of 6 supplies (including the regulator), the order might be Lower 60 (L60), U60, L20, U40, U20, and L40. Except for the regulator, which is already on, these supplies phase in simultaneously so that the current is getting a boost at regular intervals around the ring. When those supplies approach their maximum output, the next tier (which may consist of either three or six supplies) begins to turn on and fill in the gaps. The last power supplies to phase in will also be the first to phase out when the extra voltage is no longer needed.

As discussed in Chapter 2, the dipole busses have considerable inductance and will resist the change in current. The change in voltage has to be much steeper than the desired change in current. Since the inductive voltage is proportional to the change in current, \dot{P} is an indicator of how hard the supplies will have to turn on.

Remember that at the hardware level, any change in the voltage output has to be implemented by adjusting the SCR firing angles. At 8 GeV, the supplies can be on, but delaying the trigger gate may block their voltage output. By reducing the delay, more voltage is gated through and produces more current.

At line 3, PARAB, the momentum is 9.5 GeV/c and \dot{P} is 20 GeV/c per second. Clearly, things are picking up. By line 5, the beam momentum is 85 GeV/c and \dot{P} is 188 GeV/sec, but \ddot{P} has a negative value (-185 GeV/c/sec²); the rate of rise is beginning to slow down. At 2.3217 seconds, the momentum is finally at 150 GeV/c and \dot{P} is zero, in other words, flattop.

The ramp remains at flattop for nearly half a second, to 2.7292 seconds. That is sufficient time for transfer cogging (an RF process) and

Main Injector

transfer into the Tevatron. Transfer currently takes place at 2.6 seconds; the Tevatron may have already quenched by the time the ramp starts down. The cleanup abort is normally scheduled to occur sometime between 2.6 seconds and the end of flattop.

Once beam is out of the machine, during invert, the momentum rapidly plunges toward its minimum value. The trip down is faster than the trip up because without beam there is no need to maintain any sort of field quality. The magnets only have to be at their minimum value for an instant to set the hysteresis, and they are brought back to their 8.889 GeV/c level before the next cycle begins.

As mentioned earlier, during invert the SCRs are configured so that the polarity of the applied voltage is reversed, not just eliminated, so that current will not linger in the magnets. After all, we may be in a hurry to move on to the next pulse. \dot{P} assumes large negative values during this phase.

Since this chapter deals with power supplies and their regulation, the use of the tune sub page on I2 as a tool for controlling the beam will be deferred to a hypothetical chapter on “Beam Tuning.”

The sextupole and octupole tables are also kept on I2. It is important to recognize that these are not controlled by MECAR; the ramps have been placed on the same page for the convenience of the tuner.

MDAT Generation

MDAT stands for “machine data.” The data carried by MDAT is broadcast widely on the MDAT link and can be listened to by any device. MECAR, since it is already full of information about the Main Injector ramps, encodes some of that information onto MDAT. The MDAT data frame is updated at a 720 Hz rate, which is of course the same rate that the SCRs in the power supplies update. Each type of data within the frame is assigned a number, and there is a dedicated parameter for each type as well. The data generated by MECAR is based on the main dipole bus; the devices that read

Main Injector

MDAT use the information as a baseline for calculating their own ramps.

Some MDAT signals generated by MECAR include:

MDAT30: Main Injector beam momentum, in GeV/c. Remember that there is not a strict linear relationship between energy and current in the main dipoles, because the magnetic domains in the dipoles saturate at higher energies. The beam momentum is calculated from the program current based on the saturation tables stored in I2, the applications page. It is MDAT30 that the correction elements watch when playing their energy ramps (see below). If MECAR and MDAT are alive and well, this signal will be present even if there is no power to the magnets.

- MDAT31: Program Pdot. “Pdot” is the rate at which the momentum is changing. Program Pdot is a value calculated by MECAR from the program momentum. It can be useful in determining if the slew rate for a particular ramp is too high, and is used in the sextupole ramp tables.
- MDAT40: Measured current, or what the current is actually doing. It is measured by a transducer; MECAR reads the information to use in its feedback algorithms and then formats the data for broadcast on the MDAT link.
- MDAT41: Measured Idot, calculated from measured current.

MDAT uses the same fiber-optic repeater chassis as the other links. More information on MDAT in general can be found in the Controls Rookie Book.

Corrector Element Power Supplies

The correction dipoles, trim quads, skew quads, and octupoles in any given region all draw their power from a single Corrector Power Supply (CPS). The sextupoles are powered separately, as explained below. There is one CPS at each of the major service buildings, with two at MI-60. (Those who

Main Injector

remember Main Ring should consider the CPS analogous to the two bulk supplies for the old correctors. Those power supplies and magnets have been retained for use in the MI-8 line and the Recycler.) The CPS supplies, and therefore the correctors, are powered from house power and not pulsed power.

Fig. 3-6 is a representative layout of the modules at a typical service building,

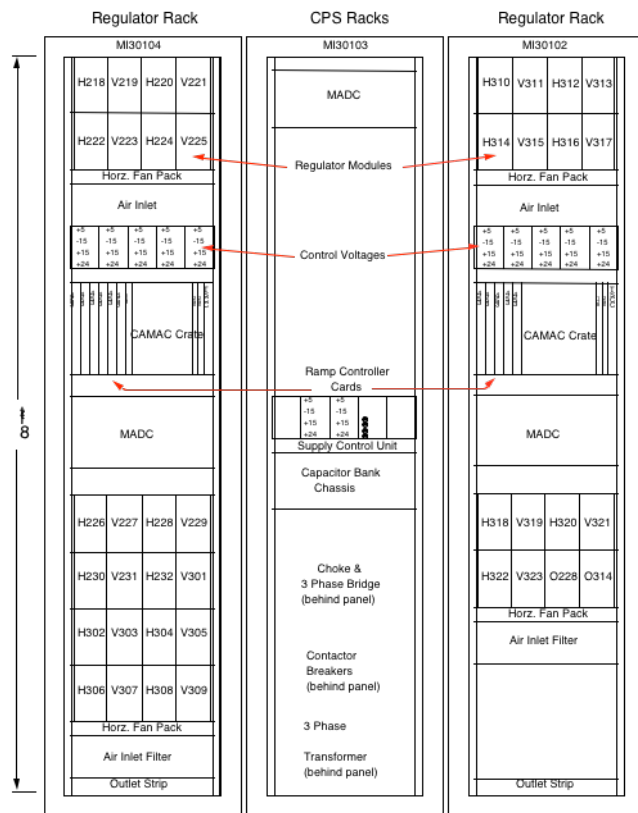


Fig. 3-6
Corrector Power Supplies at MI-30

MI-30 and Fig. 3-7 is a simplified schematic representation of the CPS. The devices at MI-30 will be used as examples repeatedly in the next few paragraphs and diagrams, not only to heighten the sense of individual drama, but also to provide continuity of the descriptions in a specific context.

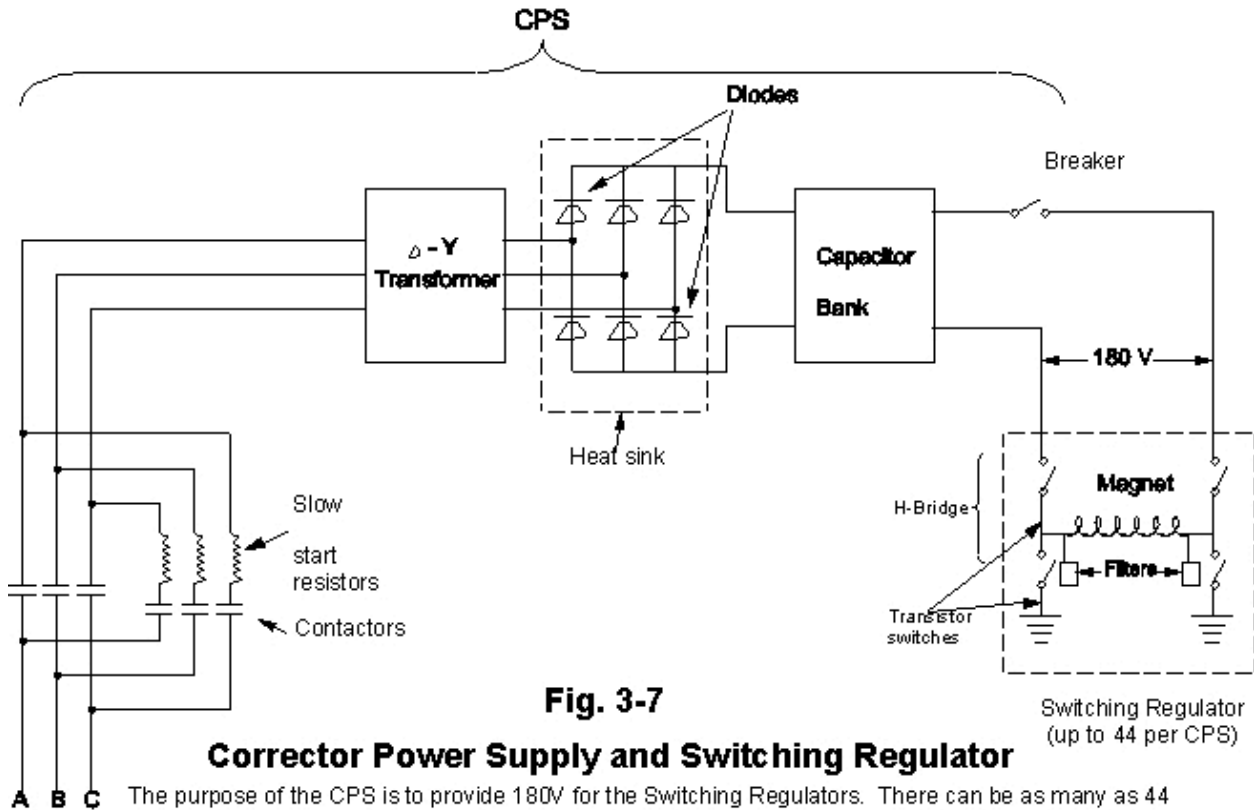
The CPS first feeds power to a number of switching regulators (only one of the switching regulators connected to the CPS is included in Fig. 3-8). Each regulator is dedicated to a single magnet or a short string of

magnets. The specific job of the CPS is to provide a constant 180V to each of the regulators, and the regulators then send the proper current to the magnets. A CPS is capable of powering up to 48 regulators, but in practice no more than 44 are actually connected.

The CPS uses 3-phase power. The three phases are first sent through a slow start circuit, that is, a small amount of current is applied at first to ensure that the phases are all matched up properly. When the circuit is energized, the slow start contactors are closed so that resistors limit the current. After a tenth of a second, the main contactors close. Since they

Main Injector

have very little resistance compared to the slow-start resistors, most of the current then passes through them and on to the supply.



The 3-phase passes through a delta-wye transformer and is rectified by a bank of diodes. (Diodes are much like SCRs in that they only allow current to pass in one direction; unlike SCRs, they are not triggered.) The diode bank needs to be kept cool, so it sits on a heat sink.

The DC current from the diode rectifier charges a capacitor bank, which in turn provides the 180V needed by the regulators.

There is a dedicated cabinet for each supply; the manual disconnect for the supply is mounted on the back door of the cabinet, to ensure that the supply is racked out before accessing the cabinet.

The CPS is air-cooled by three fans at the top of the rack. Unlike many magnet power supplies elsewhere, LCW is not used.

Main Injector

The switching regulators are the silver-colored modules found in adjacent racks. The design of a regulator is based on an H-Bridge, in which the magnet or magnet string forms the horizontal bar of the “H.” The switches—actually transistors—are found at the four poles. The switches can open and close at a 30 kHz rate, constantly adjusting the voltage applied to the load as needed. This setup allows bipolar current to flow in the magnets and updates quickly enough to keep up with the rapid changes demanded by the ramp rate.

The correction dipoles and trim quads each have a one-on-one relationship with their personal regulator. The skew quads, which as you remember from Chapter 2 are clustered into four groups of four, are assigned one regulator per cluster. The regulators assigned to the octupoles can be assigned between one and five magnets, depending on the location. The specific power supply configuration for each type of corrector is graphically displayed in a series of pictures at the end of Chapter 2.

On the front panel of each regulator is a Load Compensation Switch. The magnet loads vary from place to place, especially with the number of magnets in the string. The Load Compensation Switch anticipates what the total resistance and inductance of a string will be. The switch is not automatic; it is important to set the switch properly when replacing a regulator.

The ACNET parameter name for a given CPS is based on its location, I:CPS10 at MI-10, I:CPS20 at MI-20, etc. The parameters for the supplies at MI-60 are I:CPS60N and I:CPS60S, for “north” and “south.” The ACNET names for the regulators are named after the magnets, e.g. I:H310. If there are several magnets in series using the same regulator, the supply is named for the first magnet in the string.

Ramp Controller Cards

A ramp waveform is a series of current values, played out over the course of a machine cycle and sent to an individual regulator. The waveforms

Main Injector

for most of the correctors are stored and played from CAMAC 453 cards. The 453 card is one of a series of ramp controller cards widely used at Fermilab. The applications page I14 consolidates access to all of the ramp tables in the 400 series CAMAC cards, including those of the Main Injector, Booster, Tevatron, and the Antiproton Source.

Each 453 card is capable of storing the waveforms for up to four individual devices; each device can theoretically be assigned up to 48 waveforms, each waveform appropriate for a particular magnet.

One way in which the ramps are listed is by clock event. There can be up to 16 different clock events for each device, designated A through P. Any clock event can be assigned to any slot, but the list is usually standardized to minimize confusion. The ramp begins to play out its program when the 453 card detects the appropriate clock event, carried on TCLK. Most of the clock events used are those for Main Injector resets, e.g., a \$29 for a stacking ramp, or a \$21 for Fixed Target Tevatron injection. One clock event, a \$26 (end of beam operations), is put into Slot "L" to help smooth the transition to the beginning of the next ramp.

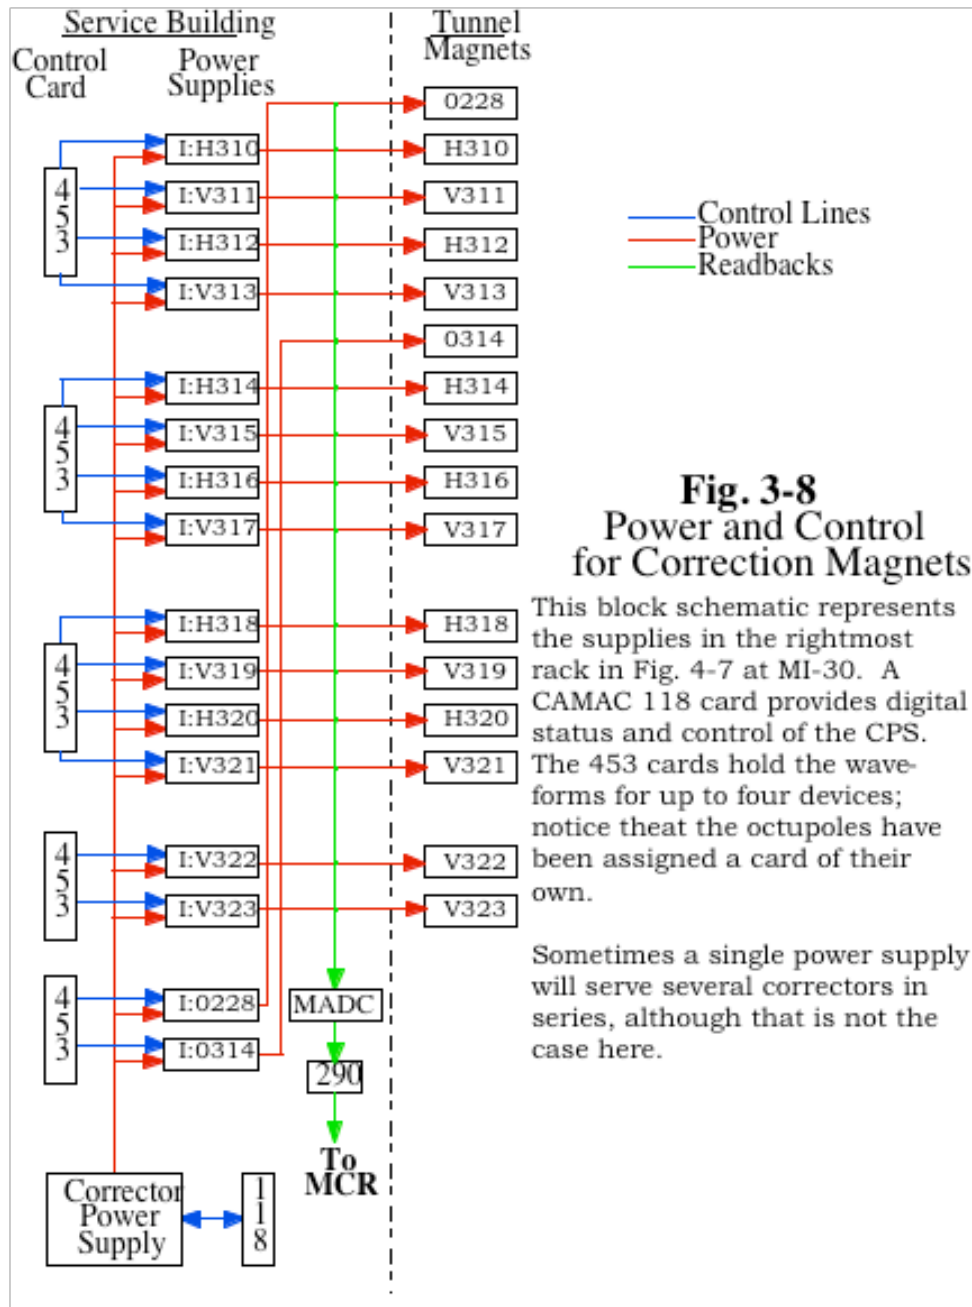
Within each clock event slot, there are three ways of defining a ramp: (1) A ramp that is a function of time, or $f(t)$; and two additional functions, $g(M1)$ and $h(M2)$, where M1 and M2 are selected MDAT signals. The MDAT signal used almost universally for correctors in the Main Injector is MDAT30, the Main Injector programmed momentum. MDAT30 is generated by the applications software of Page I2 and launched from MECAR. The $f(t)$, $g(M1)$, and $h(M2)$ ramps can be superimposed if desired.

The $f(t)$ ramp is the easiest to understand, and the least used. The ramp plays when the clock event is received, oblivious to any further instructions. In the "t" column, each slot is given a time interval. In the $f(t)$ column are the current values the power supply is expected to reach during that time interval. The card calculates the rate of change needed and sends the appropriate signal to the regulator.

Main Injector

The 453 cards use MDAT functions most of the time. Specifically, MDAT30 is assigned to the M1 function so that the regulators produce a current that is directly related to the beam energy. In that way the waveform stored in the table is still appropriate even if the ramp timing is modified.

There are applications programs, such as those for smoothing the beam orbit, whose primary function is to adjust the 453 ramps. These programs will be discussed in a mythical chapter called “Beam Tuning.”



Main Injector

Fig. 3-8 schematically shows the control lines and power distribution to the regulators in the right-most rack at MI-30. The 453 cards each talk to four power supplies, except that there aren't quite enough dipoles to go around. The last card with dipole tables only has two dipole correctors under its control. The final 453 card has two octupoles assigned to it. Although it would have been possible to put those two octupoles on the previous 453 card, the octupoles are always segregated from the other devices because they are ultimately controlled from a different source (page I2).

Often, as with the skew quadrupoles and most of the octupoles, a single regulator may drive several magnets in series. The two octupoles at MI-30, however, happen to be the only devices powered by their respective regulators (Fig. 2-24).

The digital interface to the CPS is a CAMAC 118 card. On and off commands are issued through this card, and the digital status information described below comes back to the MCR through the card.

Language Traps

Some of the corrector magnets, of course, are quadrupoles.

The 453 cards are labeled as "Quad Controllers." Here, "quad" does not mean that the cards necessarily control quadrupole magnets, but that they control four devices.

The regulators are called "Quadrant Switching Supplies." This does not mean that the regulator is necessarily powering four devices; the term refers to the four poles of the H-Bridge.

One can only imagine what electrical engineers are like on a golf course.

Main Injector

Interpreting Digital Status

When a CPS or regulator power supply trips off, the cause can usually be determined from the extended digital status (I:CPSxx) found on, for example, S53. The string “xx” refers to the house number.

Some of the status bits for the CPS are:

- **Power Supply:** This bit refers to the control voltages from the Control Unit in the middle of the CPS rack.
- **Slow Start:** This bit actually indicates that the main breakers of the CPS, which are supposed to close in a tenth of a second after the slow start breakers, actually do so. If they don't, current will continue to flow in the slow start resistors until they overheat. A Klaxon on the resistors opens and toggles the bit.
- **RMS Overcurrent:** “RMS” stands for “root-mean-square,” and is a standard method for measuring an average. If the average current produced by the CPS over the course of an entire cycle exceeds 120 amps, it will trip. The average current can be monitored through the parameter I:CPSxxI. An RMS overcurrent trip can often be fixed just by reducing the current in some of the regulators.
- **Door Interlock:** If the back door of the CPS rack is opened, the supply will trip. The door is also equipped with a manual disconnect which must be opened before entering the cabinet.
- **Cooling Fans:** The CPS is interlocked to the three fans at the top of the rack through motion sensors.
- **Transformer/Choke:** These components of the CPS are interlocked to the temperature through klaxons.
- **Reset:** Clears appropriate interlocks on the CPS, and after a short delay sends resets to the regulators. The regulators are set sequentially in groups of eight. At MI30, for example, the CPS would first reset a group beginning with H218, followed by groups beginning

Main Injector

with H226, H302, H310 and finally H318. The sequence can be watched on a string of LED's when resetting the CPS locally.

Some of the status bits for the regulator supplies (e.g. I:H310) are:

- Ramp Enabled: The ramp waveform generated by the 453 card can be enabled or disabled by toggling a bit on the 453 table page. It is also possible for the ramp to become disabled because of a failure on board the card.
- Tracking Error: A “tracking error” means that the actual current from the supply deviates from the ramp waveform playing from the 453 card. The tolerance allowed is about a quarter of an amp. Tracking errors can be diagnosed by plotting the actual current (e.g. I:H310) vs. the reference voltage (e.g. I:H310F). (In the Main Ring the 453 cards were updated at a 15 Hz rate, and tracking errors were common. The cards have now been upgraded to update at 720 Hz, which is expected to alleviate the problem.)
- Analog alarms are impractical for corrector supplies, because the amount of current can legitimately be anywhere within a large range. However, a digital tracking error will post an alarm if a regulator is not producing current or has other serious problems following the ramp program.
- DC Overcurrent: The regulator supplies are capable of producing up to ± 20 amps. A current limit of about 19 amps is set in the hardware that will trip the supply if exceeded. This bit is keyed to the instantaneous value of the current rather than to an average. The problem can often be diagnosed by plotting the output of the supply to see if it crosses over the limit at any point.
- Ground Fault: Obvious.
- RMS Overcurrent: This is the same as for the CPS, an average current over an entire cycle. The limit for an individual regulator is 12 amps in

Main Injector

the case of the dipoles. As with the CPS, an RMS overcurrent for an individual regulator supply can often be cleared by reducing the overall current that it produces over a cycle.

- Low Input Voltage: This means that the regulator is not seeing enough voltage from the CPS, which it reads from the capacitor bank. (When turning off the CPS locally, the LED for the input voltage on the regulator supplies can be seen lingering for a few seconds as the capacitors bleed down.)
- External Permit: This is the permit issued by the CPS to the blocks of 8 regulators, causing them to turn on sequentially (see under CPS reset).
- Overtemp: This bit looks specifically at the filter resistors and heat sink inside the regulator chassis. Temperature sensors are used here instead of klixons.

Ramped Sextupoles

All of the horizontal sextupoles are in series with each other; the vertical sextupoles form a similar but independent loop (Figs. 2-22, 2-23). Unlike the other correctors, the ramped sextupoles only require one power supply per loop. The large power supplies, I:SEXHPS and I:SEXVPS, are located in the MI-52 Service Building. Power to the sextupole supplies comes from Feeders 96 and 97, so power is removed when MOS 89 is opened.

There are numerous sextupole magnets for each plane, so connecting one terminal of the power supply to every other magnet, and catching the other half with the return path balances the inductive load. For example, the positive lead of I:SEXHPS is connected to S514, but the cable from S514 skips over S512 and is instead connected to S510. The loop leapfrogs nearly all the way around the ring before reaching S614, where it turns around and connects to all of the horizontal sextupole magnets that were missed the first time.

Main Injector

Digital control of the power supplies is through PLCs, which in turn are controlled through the VME crate in the electronics room of the service building.

The sextupole ramp waveforms, I:SEXH and I:SEXV, are sent to the power supplies from a CAMAC 453 card in Crate \$5B. The waveforms are not only a function of MDAT30 (Main Injector momentum), but also of MDAT31 (Pdot, or the rate of change of momentum).

As you may recall from Chapter 2, ramped sextupoles are used to control the chromaticity of the ring. For tuning convenience, the chromaticity tables are normally adjusted from I2, but that should not imply that they are directly controlled by MECAR; they are not. There are separate chromaticity tables for each type of Main Injector reset.